April 2005 Report of the Tevatron BPM Upgrade
wbs item 1.4.5.4 of the Run 2 Luminosity Upgrade Project
Bob Webber, Stephen Wolbers
May 5, 2005

Project Manager's Summary:

Great progress was made in April on the TeV BPM Upgrade. Measurement errors, characterized as turn-by-turn (TBT) and TCLK problems, were solved or vastly reduced in frequency, leading to confidence in the system performance and a go-ahead for full installation. By the end of the month 23 out of the 27 houses were installed and connected and were in use by the Tevatron.

The TBT and missing TCLK problems became a serious concern in early April when it became clear that there was an effect causing data to be collected incorrectly in some fraction (10-15%) of the injection TBT. The project dedicated a large effort to understanding the problem, recreating it on the teststand and ultimately solving it with changes to the timing board firmware.

Work continued on assembling and testing crates for installation during April. By the end of the month all 27 crates had been assembled and tested and were ready for installation. This was a major accomplishment for the project.

Final software and hardware tasks will be the focus for the project during the next month. The project also will be working on the transition to operations.

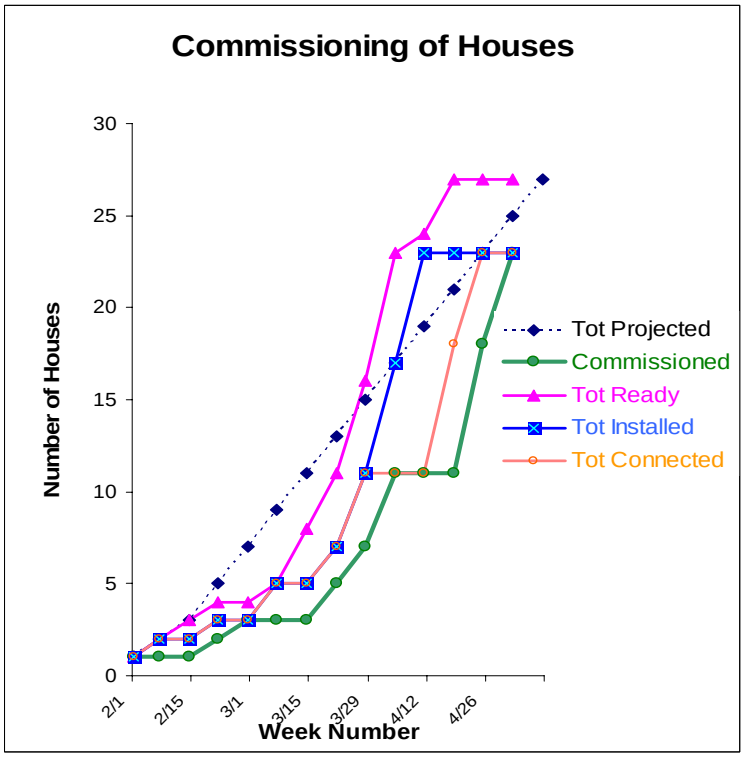

Resources Used in April 2005:

The total number of FTE-months devoted to the project in calendar April 2005 from the Computing Division was reported to be 7.1 FTE-months with 22 people contributing. The total number of FTE-months devoted to the project from the Accelerator Division was 2.0 FTE-months with 9 people contributing. The total effort from both Divisions was 9.1 FTE-months. The following table gives the estimated or reported effort for both divisions (in FTE-months) since August of 2003.

Month	AD Effort	CD Effort	Total Effort
August, 2003	1.2	2.3	3.5
September, 2003	1.4	4.1	5.5
October, 2003	5.4	6.0	11.4
November, 2003	1.6	5.0	6.6
December, 2003	1.4	4.4	5.8
January, 2004	1.7	5.1	6.8
February, 2004	2.3	6.7	9.0
March, 2004	2.1	7.6	9.7
April, 2004	2.0	7.7	9.7
May, 2004	1.4	8.3	9.7
June, 2004	1.6	8.7	10.3
July, 2004	2.0	8.1	10.1
August, 2004	1.5	8.0	9.5
September, 2004	2.3	8.4	10.7
October, 2004	1.1	10.5	11.6
November, 2004	1.8	9.9	11.7
December, 2004	1.1	7.5	8.6
January, 2005	1.1	10.0	11.1
February, 2005	2.5	10.9	13.4
March, 2005	1.3	8.6	9.9
April, 2005	2.0	7.1	9.1
SUM (through Apr, 2005)	38.8	154.9	193.7

The effort is consistent with the project finishing certain testing tasks and moving into final installation and commissioning. The effort listed here is time worked and does not include vacation, sick leave, holidays, etc.

Purchase requisitions/procard obligations in April, 2005:

No purchases in April.

Milestones:

The March milestones will not be achieved due to the delay in installation and commissioning. All milestones will be achieved by the end of May, 2005. The old and new proposed dates are:

1.4.5.4.4.1.7 All Tev BPM crates functionally available (except for F bldg.) or Installed:

Old date: 3/4/05 New date: 5/16/05

1.4.5.4.5.1.4 Tev BPM Electronics commissioning complete

Old date: 3/30/05 New date: 5/16/05

1.4.5.4.6.8 Tev BPM Upgrade Operational

Old date: 3/30/05 New date: 5/31/05

Meetings held, Reports Given:

Meetings were held in April on the following dates:

Project Meetings: April 4,6,7,11,13,18,20,27

Documents:

The following documents were written and added to the Accelerator Division Document Database in April:

792-v12 Minutes from the Tevatron BPM Upgrade Project Meetings Steve Wolbers *et. al.* 28 Apr 2005

860-v32 Tevatron BPM Software Specifications Jim Steimel *et. al.* 26 Apr 2005

1787-v1 Checking the Abort Buffers Robert K Kutschke 21 Apr 2005

907-v2 Monthly Reports of the Tevatron BPM Upgrade Project Steve Wolbers 10 Apr 2005

Hand Triggered TBT Data With the Upgraded Tevatron BPMS Robert K Kutschke 06 Apr 2005

Subproject Leader Reports:

Technical Coordinator: Jim Steimel

During the month of April, most of the systems were installed in the Tevatron service buildings. We held off on commissioning many of the systems until the issues related to the turn-by-turn measurements were resolved. Enough modifications to the front-end software were made to reduce the statistics of the error, so that Tevatron operations would not be handicapped. We then commissioned systems at a very fast rate. All but four of 27 houses are currently installed and commissioned. The majority of the coordination effort involved verifying proper timing for all the systems and running system diagnostics to check for unconnected or misconnected cables.

A solution to the race condition between the Echotek clock and sync inputs was discovered. This was tested positively for two houses, and half of the BPMs were corrected by the end of the month.

There was an investigation into the possibility of using the change in phase difference between signals from the different plates as a secondary way of measuring beam position. This could then be used to verify scalers and electrical offsets. However, studying data revealed that the linear dependence of position to phase difference varies considerably from BPM to BPM. This is due to the mismatch in impedance at the opposite end of the pickup. The reflections will have a large effect on the relative phase difference between the signals measured from each plate.

Effort in May will see the completion of installation and commissioning of the BPM hardware. The race conditions will be eliminated, and work will begin on an offline diagnostics application.

Hardware Tracking/Testing: Tim Kasza

During the month of April, our contributions have mainly been focused on helping out with the on-going work in the Tevatron service buildings. This has involved a variety of work which included, installing 36 foot clock cables, repairing LEMO connectors on clock cables, connecting proton cables on previously installed crates, replacing prototype filter boards, swapping out malfunctioning Echotek cards, installing termination boards and the continued installation of new crates. In order to free up crate CPUs for the remaining crate installations, the overall total number of TeV-BPM test stands on FFC 3 were consolidated down from 7 to 4. So far 23 out of 27 crates have been installed, of which, the remaining 4 crates (A0, F3, F4, and F1) have been assembled, tested and now awaiting an installation opportunity.

Some effort in other areas includes the repair of 7 Filter boards which were previously damaged during the debugging of the diagnostic signal problem. Additionally, we processed and shipped 6 defective filter pairs back to Lark for warranty replacement. Several Echotek boards have returned from warranty repair and are in the process of being tested to see if they will pass our tests.

Front-end/DAQ software: Margaret Votava

Most of the month of April was spent on finding and fixing the 6 thousand turn problem first noticed in the month of March, after the resolution of the outliers in turn-by-turn data. Additionally there were discussions about the safe mode implementation, usage of the user TCLK buffer to gather data for calculation of pbar cancellation coefficients and small bug fixes.

The six thousand problem was first noticed in injection turn-by-turn data. The first set of points had no valid data, as if the beam arrived at turn six thousand. A similar problem could be seen for hand-triggered measurements, whose data presented a discontinuity in position starting at different points for every measurement.

The test stand was configured to reproduce the problem, but there was no success. The detection required an algorithm that could detect very small jumps in the position information or verification of a discontinuity in the phase pattern. The algorithm previously used in the outliers problem could not identify the problem.

Discussions at the regular BPM meetings suggested some hypotheses: Echotek boards are not writing to the correct memory location; There is an addressing problem on the DMA; or the timing board is triggering prematurely.

As a first measurement the Echotek driver code was modified to clear the FIFO memory for the first nine thousand entries and write a flag at the location corresponding to eight

thousandth entry. The flag should never be overwritten since the Echotek boards are not supposed to take more than eight thousand samples.

The injections following the installation of the new driver did not show signs of the six thousand turn problem. However, it was verified that hand-triggered turn-by-turn measurements were still affected. Furthermore, the flag in the Echotek FIFO was overwritten. Similar experiments on the test stand verified that the flag was been overwritten after some number of turn-by-turn measurements. The number of elements written to the FIFO varied from board to board, where the boards located farther from the processor had fewer extra elements written.

In order to debug the problem some modifications were introduced to the timing board firmware, allowing the number of sync triggers sent out to be counted. It was verified that when the FIFO had more elements than expected the sync counters had values higher than 8K, indicating extra syncs produced. The logging messages in the software showed that when the problem is detected, there is an inversion of the interrupts received by the processor. That indicated a potential problem in the timing board firmware. The following paragraphs explain in more details the problems seen.

For a turn-by-turn or injection measurement the software first configures the timing board and then it reloads the Echotek setup. After the timing board is configured it is expected to send 8192 triggers to the Echoteks after it receives a BSYNC 0x7C (injection) or 0xDA (for hand-triggered TBT) start event.

The steps to prepare the timing board for a turn-by-turn measurement include the change of the gate count from 1 (closed orbit) to 8192.

A race condition in the timing board firmware makes it start a measurement immediately after it is configured for a turn-by-turn. It triggers as many times as indicated in the gate count register (1 for closed orbit or 8192 for TBT).

Even when the early triggering happens the timing board starts sending syncs out again (8192 of them) when the start event (BSYNC 0x7C or 0xDA) is received.

The effect of early triggering depends on the sequence of events in software. In short it takes the following steps when switching from closed orbit to turn-by-turn:

1) set timing board

2) load Echotek configuration

3) set number of trigger to be received by the Echoteks

4) Echoteks are armed

Step 2 takes 41ms, and during that time the Echoteks are not armed. They are armed only after the configuration is loaded. The arming order starts from the board closer to the processor.

When the timing board triggers at the expected time (after receiving BSYNC 0x7C or 0xDA) the Echoteks are ready to take data. For injection there are at least 2.6 seconds between the Echoteks arm and the 0x7C event. For hand-triggered there is a variable delay up to 1 second (0xDA is usually received at 1Hz rate). If the 0xDA is received before the Echoteks are armed then they will collect 8192 - X samples and at the next 0xDA they will collect the remaining X samples (This explains the hand-triggered failure at A3 with the memory clearing enabled).

If the timing board starts triggering early during an injection, the Echoteks will not take any samples for the first triggers issued while it is reloaded. That takes approximately 41ms (1975 turns). After the boards are armed they immediately start taking data, and will do so for the next 6217 turns. There won't be any valid data on these turns because there is no beam at that time. Around 2.5 seconds later the timing board receives the BSYNC 0x7C event, which starts the triggering again. However at this point the Echoteks have already collected 6217 turns, so they will collect the remaining 1975 and stop. At that point the data is collected by the software and we see that the beam starts around turn 6000.

In order to debug the problem a modification in the driver was introduced: the first 32KBytes of the Echotek FIFOs were zeroed and a flag written to the location equivalent to 8K samples. With this modification the time spent reloading the Echotek setup jumped from 41ms to 250ms. Now even when the timing board triggered early the Echoteks would ignore it because they were not armed. At the next start event the Echoteks collect the 8192 points at the correct time, and the software reads them out.

The injection data, when the 6000 turns problem happened, showed that boards get the first beam at different turns. This is likely to be caused by the arming order. There were tests to change the order by modifying the Echotek setup load order and swapping the cables between boards, but that didn't have an effect because the arming order remained the same.

Besides the problem described above, during April there were discussions about pbar cancelation and possible use of the user TCLK buffer to get closed orbit measurements. This requires the configuration of a TCLK that is triggered before and after the helix opens. Rob Kutschke also would like to get measurements between proton injections. The total number of points is about 35 to 40, and the user buffer has 128 points available.

Online software: Brian Hendricks

During the past month, the main accomplishment was that the remaining database scripts to convert the BPM configuration database were completed. Bugs were fixed in the BPM library and the generic BPM display application with regards to reading and plotting BPM intensity values. The BPM diagnostic program added the display of software version numbers from the front ends. The phase plot was also modified to improve readability.

Offline software: Rob Kutschke

In April my main work was to test the data quality in the installed houses. The main part of this effort was to characterize and troubleshoot the so called "6000 turn problem" that was first observed in late March. This work also uncovered many small problems which were promptly fixed by others in the team. These problems included loose or bad cables, bad boards, formatting problems in the data produced by the front ends, new ways for the front end computers to get stuck waiting forever and so on.

At the start of April, the 6000 turn problem was observed only in injection turn by turn data. Moreover it was intermittent. This provided few opportunites to characterize the problem and test candidate solutions. As described in Beams-doc-1775, I developed a way to reproduce the problem in triggered turn by turn data with steady state beam conditions. This allowed the problem to be transferred to the test stand where it was eventually solved by others in the group.

I certified that when the 6000 turn problem occurs on injection TBT data, the first point above threshold really is the point from the first turn. This allowed Luciano to install a work-around to supply good first turn data for use by the operators. This work-around will be removed once we have verified that recent fixes have indeed solved the problem.

For part of the month, 12 houses were installed but only the pbar cables were connected. I developed ways to check these houses using the proton contamination on the pbar cables. This allowed us to debug a number of problems before the proton cables were available. When the proton cables were connected, these houses worked immediately.

I worked with Roger Tokarek to help him developed a console application, W133, to acquire all of the injection turn by turn data from all of the new front ends. This is now working and in regular use. We expect a few improvements before that application is complete.

I transferred the technology for using the injection turn by turn data to members of the Tevatron department, Yuri Alexahin, Eliana Gianfelice-Wendt, and Vahid Ranjbar. They are now using this data to measure properties of the Tevatron lattice and other projects. They also have access to my saved data sets of injection turn by turn data going back several months.

Following a quench on April 21, I inspected the fast and slow abort buffers to verify that they held correct and useful information. This is described in Beams-doc-1787. One can see the beginnings of the loss of beam in the last few turns of the fast abort buffer.

Mike Martens, Luciano Piccoli, Brian Hendricks and I developed a plan to acquire the data for the calibration of the cancellation of the proton contamination on the pbar cables. The plan is to develop a buffer similar to the existing profile buffer, triggered with a new TCLK event. Mike later discovered that we will not be able to deploy a new TCLK event until the fall. Mike then developed a new plan which can be used in the interim. This will be developed and deployed in the next weeks.

Marc Mengel has designed the configuration/calibration database and has loaded the existing information into it.